



AFRL-RI-RS-TR-2017-230

PRiFi NETWORKING FOR TRACKING-RESISTANT MOBILE COMPUTING

YALE UNIVERSITY

NOVEMBER 2017

FINAL TECHNICAL REPORT

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED

STINFO COPY

**AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE**

NOTICE AND SIGNATURE PAGE

Using Government drawings, specifications, or other data included in this document for any purpose other than Government procurement does not in any way obligate the U.S. Government. The fact that the Government formulated or supplied the drawings, specifications, or other data does not license the holder or any other person or corporation; or convey any rights or permission to manufacture, use, or sell any patented invention that may relate to them.

This report is the result of contracted fundamental research deemed exempt from public affairs security and policy review in accordance with SAF/AQR memorandum dated 10 Dec 08 and AFRL/CA policy clarification memorandum dated 16 Jan 09. This report is available to the general public, including foreign nations. Copies may be obtained from the Defense Technical Information Center (DTIC) (<http://www.dtic.mil>).

AFRL-RI-RS-TR-2017-230 HAS BEEN REVIEWED AND IS APPROVED FOR PUBLICATION IN ACCORDANCE WITH ASSIGNED DISTRIBUTION STATEMENT.

FOR THE CHIEF ENGINEER:

/ S /

ROBERT L. KAMINSKI
Work Unit Manager

/ S /

WARREN H. DEBANY JR.
Technical Advisor, Information
Exploitation and Operations Division
Information Directorate

This report is published in the interest of scientific and technical information exchange, and its publication does not constitute the Government's approval or disapproval of its ideas or findings.

REPORT DOCUMENTATION PAGE

**Form Approved
OMB No. 0704-0188**

The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE (DD-MM-YYYY) NOV 2017			2. REPORT TYPE FINAL TECHNICAL REPORT		3. DATES COVERED (From - To) FEB 2016 – MAY 2017	
4. TITLE AND SUBTITLE PRiFi NETWORKING FOR TRACKING-RESISTANT MOBILE COMPUTING			5a. CONTRACT NUMBER FA8750-16-2-0034			
			5b. GRANT NUMBER N/A			
			5c. PROGRAM ELEMENT NUMBER			
6. AUTHOR(S) Joan Feigenbaum			5d. PROJECT NUMBER DHS1			
			5e. TASK NUMBER 4Y			
			5f. WORK UNIT NUMBER AL			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Yale University 51 Prospect St New Haven CT 06511-8937				8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Research Laboratory/RIG 525 Brooks Road Rome NY 13441-4505				10. SPONSOR/MONITOR'S ACRONYM(S) AFRL/RI		
				11. SPONSOR/MONITOR'S REPORT NUMBER AFRL-RI-RS-TR-2017-230		
12. DISTRIBUTION AVAILABILITY STATEMENT Approved for Public Release; Distribution Unlimited. This report is the result of contracted fundamental research deemed exempt from public affairs security and policy review in accordance with SAF/AQR memorandum dated 10 Dec 08 and AFRL/CA policy clarification memorandum dated 16 Jan 09						
13. SUPPLEMENTARY NOTES						
14. ABSTRACT As the most serious cyber-attack threats rapidly shift from untargeted toward increasingly targeted methods, it is becoming correspondingly more crucial for organizations to protect the identity and location-privacy of their members against malicious tracking and surveillance. We propose to develop PriFi, an anti-tracking and location-private network access mechanism to protect members of an organization both while on-site (via privacy-protected WiFi networking) and while off-site (via privacy-protected Virtual Private Networking or VPN).						
15. SUBJECT TERMS Anonymity; Location privacy; Tracking resistance						
16. SECURITY CLASSIFICATION OF:		17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 15	19a. NAME OF RESPONSIBLE PERSON ROBERT L. KAMINSKI		
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U	19b. TELEPHONE NUMBER (Include area code)			

TABLE OF CONTENTS

Section	Page
List of Figures	ii
1.0 Summary.....	1
2.0 Introduction.....	1
3.0 Methods, Assumptions, and ProcedureS	2
4.0 RESULTS AND DISCUSSION.....	2
4.1. PriFi: A Low-Latency and Tracking-Resistant Protocol for Local-Area Anonymous Communication.....	2
4.2. Analysis of Scheduling Algorithms for PriFi	3
4.3. Analysis of the PriFi Protocol.....	3
4.4. Avoiding The Man on the Wire: Improving Tor's Security with Trust-Aware Path Selection.....	4
4.5. Scalable Bias-Resistant Distributed Randomness	5
4.6. Privacy-Preserving Lawful Contact Chaining.....	5
5.0 Conclusions.....	7
6.0 References.....	8
7.0 APPENDIX AND BIBLIOGRAPHY	9
LIST OF ACRONYMS	10

LIST OF FIGURES

Figure		Page
Figure 1: Basic PriFi Setup		3
Figure 2: What We Have: A Cloud of Secret Mass Surveillance Processes		6
Figure 3: What We Require: Open Warrant-Based Processes for Lawful Electronic Surveillance, Creating a “Privacy Firewall”		7

1.0 SUMMARY

As the most serious cyber-attack threats rapidly shift from untargeted toward increasingly targeted methods, it is becoming correspondingly more crucial for organizations to protect the identity and location-privacy of their members against malicious tracking and surveillance. We propose to develop PriFi, an anti-tracking and location-private network access mechanism to protect members of an organization both while on-site (via privacy-protected WiFi networking) and while off-site (via privacy-protected Virtual Private Networking [VPN]).

PriFi builds on state-of-the-art accountable-anonymity technology from the Dissent project at Yale and UT Austin, a major effort in tracking-resistant communication funded by DARPA under the SAFERWarfighters program. The Dissent project created and developed proof-of-concept prototypes of the first practical tracking-resistance technology based on the information-theoretic Dining Cryptographers or DC-nets approach to anonymous communication. PriFi continues earlier efforts with a concentration on performance and usability enhancements by formulating a novel remote-trustees security model that separates trust from the critical performance loop without weakening security.

2.0 INTRODUCTION

PriFi enables mobile-device users to communicate as usual via voice and video calls, text messages, or Web browsing, but it simultaneously enables them to make the network-access activities of hundreds or thousands of users “blend together” into an amorphous cloud, cryptographically impenetrable by *external or internal* adversaries attempting to intercept, track, or eavesdrop on individuals.

PriFi is aimed at intra-organizational, tracking-resistant communication, but multiple collaborating organizations may also use it by joining together into a bigger “cloud” that is treated by PriFi as one organization. Our long-term goal is a hardened, fully deployable PriFi implementation that can be used in conjunction with other relevant services, including Tor for tracking-resistant inter-organizational communication, distributed protocols for the generation of high-quality shared random bits, and privacy-preserving surveillance and law-enforcement protocols.

During the 15.5 months of this project, we first did a proof-of-concept implementation of PriFi and then created a substantially improved, more efficient, and more usable PriFi version 2. We integrated version 2 with existing authentication systems and conducted numerous performance tests, including distributed tests running simultaneously at EPFL and Yale. We also worked on closely related problems that will affect the long-term deployability of PriFi, including trust-aware path selection in Tor, distributed randomness, and surveillance infrastructure that preserves the privacy of innocent users.

Four published papers, one Masters Thesis, and an as-yet-unpublished analysis of PriFi were supported in whole or in part by this award. They are listed in the Bibliography below.

3.0 METHODS, ASSUMPTIONS, AND PROCEDURES

We used standard experimental methodology and procedures throughout this project. The only assumption that we made is that the cryptographic building blocks are unbreakable by realistically resourced adversaries; that assumption is necessary only for the security analysis – not for the experimental performance analysis.

4.0 RESULTS AND DISCUSSION

4.1. PriFi: A Low-Latency and Tracking-Resistant Protocol for Local-Area Anonymous Communication

Popular anonymity mechanisms such as Tor [1] provide low-communication latency but are vulnerable to traffic-analysis attacks that can de-anonymize users. Moreover, known traffic-analysis-resistant techniques such as Dissent [2] are impractical for use in latency-sensitive settings such as wireless networks. We propose PriFi, a low-latency protocol for anonymous communication in local-area networks that is provably secure against traffic-analysis attacks. This allows members of an organization to access the Internet anonymously while they are on-site, via privacy-preserving WiFi networking, or off-site, via privacy-preserving virtual private networking (VPN).

PriFi reduces communication latency using a novel client/relay/server architecture in which a set of servers computes cryptographic material in parallel with the clients to minimize unnecessary communication latency. We also propose a technique for protecting against equivocation attacks, with which a malicious relay might de-anonymize clients. This is achieved without adding extra latency by encrypting client messages based on the history of all messages they have received so far. As a result, any equivocation attempt makes the communication unintelligible, preserving clients' anonymity while holding the servers accountable.

Figure 1 depicts the basic PriFi setup.

For a more detailed account, please refer to [3].

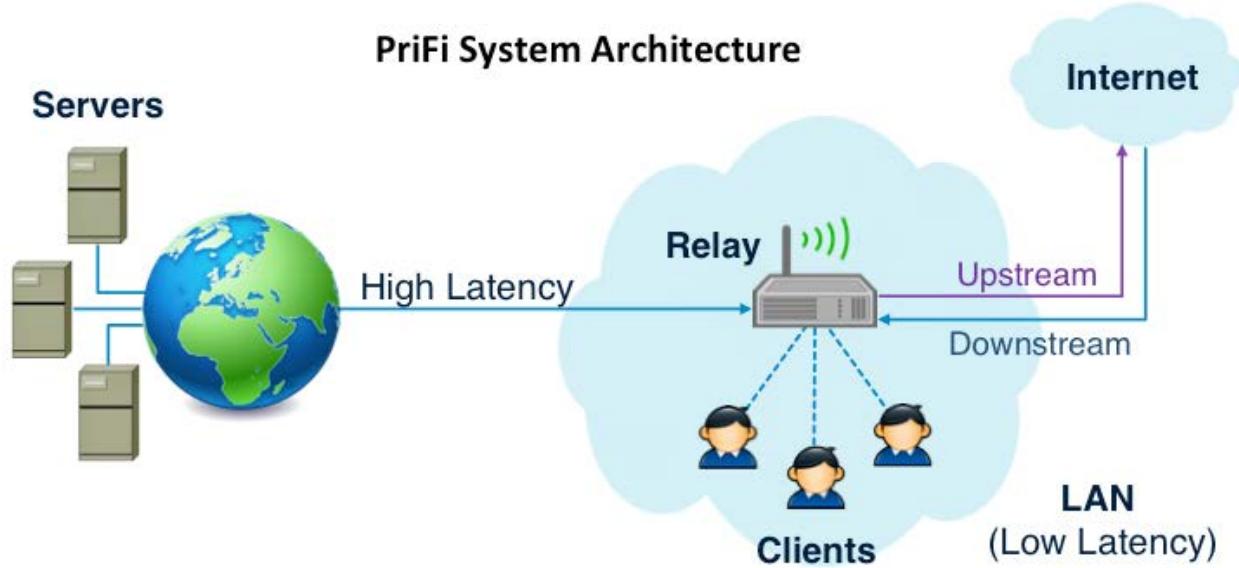


Figure 1: Basic PriFi Setup

4.2. Analysis of Scheduling Algorithms for PriFi

PriFi can incur significant extra cost if resources are scheduled carelessly. We analyze three approaches to resource scheduling in the context of anonymous communication using DC-nets: contention-based scheduling, reservation maps, and fixed scheduling. We determine which is best suited to particular scenarios by first analyzing scheduling algorithms theoretically and then running simulations on real datasets.

After running our simulator on multiple datasets with multiple strategies, we reach the same conclusion as we did in our theoretical analysis. Reservation maps are best suited in most cases, because they are able to handle traffic change more easily by letting users reserve slots. Fixed scheduling is a viable alternative when the network activity is high and when the vast majority of users are active all the time. To deal with long period of inactivity, adding mechanisms to slow down the system is a good way to save resources.

For a more detailed account, please refer to [4].

4.3. Analysis of the PriFi Protocol

We give a fuller specification and analysis of the PriFi protocol than can be found in [3]. In particular, we explain in detail the system model, threat models, and goals of the protocol.

Our system model comprises a set of n clients (or users) who want to access the Internet anonymously. They do so through a local-area network, where they connect to a relay (in practice, probably a router) that can process normal TCP/IP traffic as well as PriFi traffic (see Figure 1). Outside the local-area network, there is a set of m servers the role of which is to assist the relay in the anonymization process. These servers may be distributed globally in order to maximize diversity and collective trustworthiness. Therefore, *the connection between servers*

and the relay is high-latency. We define *downstream communication* as the data from the Internet to one of the clients and *upstream communication* as the data from one of the clients to the Internet.

The PriFi threat model assumes that an adversary controls the relay, up to $n-2$ of the clients, and up to $m-1$ of the servers. A node (client, server, or relay) is said to be *honest* if it follows the protocol faithfully and does not collude with or leak sensitive information to any other node. A *dishonest* (or *malicious*) node is controlled by the adversary and hence may deviate from the protocol or leak sensitive information. The servers satisfy the requirements of the *anytrust model* [2]: At least one server is honest, *but they are all highly available*, and the clients need not know or choose which server to trust. In particular, malicious nodes, in addition to deviating arbitrarily from the protocol, can collude with each other to de-anonymize honest clients and disrupt the communication.

The goal of PriFi is to enable members of an organization to communicate anonymously; more precisely, when using PriFi, no one inside or outside of the organization can track the communication or attribute individual messages to their senders and receivers. From the point of view of an off-site member who is using the organization’s PriFi network to communicate remotely with other members or with an arbitrary site on the Internet, PriFi is similar to a low-latency VPN service: It receives data from and sends data to the applications running on the user’s computer. The relay acts as the other end of the VPN, relaying data between the Internet and the clients. However, unlike traditional VPN services, the relay is not trusted; it may maliciously (possibly by colluding with other untrusted entities) attempt to de-anonymize the clients. The anytrust group of servers collectively allows PriFi to protect the clients from de-anonymization by the VPN service without adding latency to the critical communication path.

We show that PriFi meets these goals by showing that it provides *equivocation protection* and *disruption protection*. The former means that the protocol will expose a malicious relay that tries to de-anonymize by sending different downstream messages to different clients. The latter means that, if one or more participants try to actively disrupt the protocol, at least one of them will be de-anonymized and exposed to all other participants, who can then exclude the malicious actor from participating in subsequent protocol executions.

For a more detailed account, please refer to [5].

4.4. Avoiding The Man on the Wire: Improving Tor’s Security with Trust-Aware Path Selection

As explained in Section 2, our long-term goal is a hardened, fully deployable PriFi implementation that can be used in conjunction with other relevant services, including Tor for tracking-resistant inter-organizational communication. Therefore, we examine known vulnerabilities in Tor and explore ways to mitigate them.

Tor users are vulnerable to de-anonymization by an adversary that can observe some Tor relays or some parts of the network. We demonstrate that previous network-aware path-selection algorithms that propose to solve this problem are vulnerable to attacks across multiple Tor connections. We propose that users use trust to choose the paths through Tor that are less likely to be observed, where trust is flexibly modeled as a probability distribution on the location of the user’s adversaries. We present the *Trust-Aware Path Selection algorithm* for Tor that helps users avoid traffic-analysis attacks while still choosing paths that could have been selected by many other users. We evaluate this algorithm in two settings using a high-level map of Internet routing:

(i) users try to avoid a single global adversary that has an independent chance to control each Autonomous System organization, Internet Exchange Point organization, and Tor relay family, and (ii) users try to avoid de-anonymization by any single country. We also examine the performance of Trust-Aware Path selection using the Shadow network simulator [6].

For a more detailed account, please refer to [7].

4.5. Scalable Bias-Resistant Distributed Randomness

Also relevant to long-term, large-scale deployability of PriFi is access to high-quality, public randomness. Indeed, bias-resistant, public randomness is a critical component in many distributed protocols, in the anonymous-communication domain and almost all others. Generating public randomness is hard, however, because active adversaries may behave dishonestly to bias public random choices toward their advantage. Existing solutions do not scale to hundreds or thousands of participants, as is needed in many decentralized systems.

We propose two large-scale distributed protocols, *RandHound* and *RandHerd*, which provide publicly verifiable, unpredictable, and unbiased randomness against Byzantine adversaries. RandHound relies on an untrusted client to divide a set of randomness servers into groups for scalability, and it depends on the pigeon-hole principle to ensure output integrity, even for non-random, adversarial group choices. RandHerd implements an efficient, decentralized randomness beacon. RandHerd is structurally similar to a BFT protocol, but uses RandHound in a one-time setup to arrange participants into verifiably unbiased random secret-sharing groups, which then repeatedly produce random output at predefined intervals.

Our prototype demonstrates that RandHound and RandHerd achieve good performance across hundreds of participants while retaining a low failure probability by properly selecting protocol parameters, such as a group size and secret-sharing threshold. For example, when sharding 512 nodes into groups of 32, our experiments show that RandHound can produce fresh random output after 240 seconds. RandHerd, after a setup phase of 260 seconds, is able to generate fresh random output in intervals of approximately 6 seconds. For this configuration, both protocols operate at a failure probability of at most 0.08% against a Byzantine adversary.

For a more detailed account, please refer to [8].

4.6. Privacy-Preserving Lawful Contact Chaining

Tracking resistance as provided by PriFi is one component of the very broad and, to some extent, amorphous notion of *online privacy*. In recent years, a great deal of policy debate has assumed, either implicitly or explicitly, that there is a need for “balance” between online privacy and national security. Both sides of these “balance” arguments presume that security and privacy represent a zero-sum tradeoff, a presumption that we believe is false – not just on public-policy grounds but also for technical reasons. We believe that, by deploying appropriate technology in the context of sound policy and the rule of law, citizens can have both strong national security and strong online-privacy protections.

We first addressed this issue in [9], before the PriFi project started. Consistent with both US Constitutional and human-rights principles that allow government “search and seizure” in private spaces only via warrant processes grounded in public law, we proposed that any electronic-surveillance activity searching or otherwise touching private user data or metadata

must likewise be implemented via *open, public* processes that protect the privacy of innocent, untargeted users (*i.e.*, users who are not the subjects of legitimate warrants). We formulated an openness principle comprising two main planks: (1) Any surveillance or law-enforcement process that obtains or uses private information about untargeted users shall be an open, public, unclassified process, and (2) Any secret surveillance or law-enforcement processes shall use only: (a) public information and

(b) private information about targeted users obtained under authorized warrants via open surveillance processes. Instead of the deplorable situation that we have now, in which secret processes govern the bulk collection of private data about all Internet users (depicted in Figure 2), we propose the creation of a “privacy firewall” in which open processes ensure the protection of private information about innocent parties (depicted in Figure 3). As a concrete case study, we designed, prototyped, and tested in [9] a metadata-query system based on a mature and practical privacy-preserving protocol for set intersection, an operation that law-enforcement and intelligence agencies have used effectively.

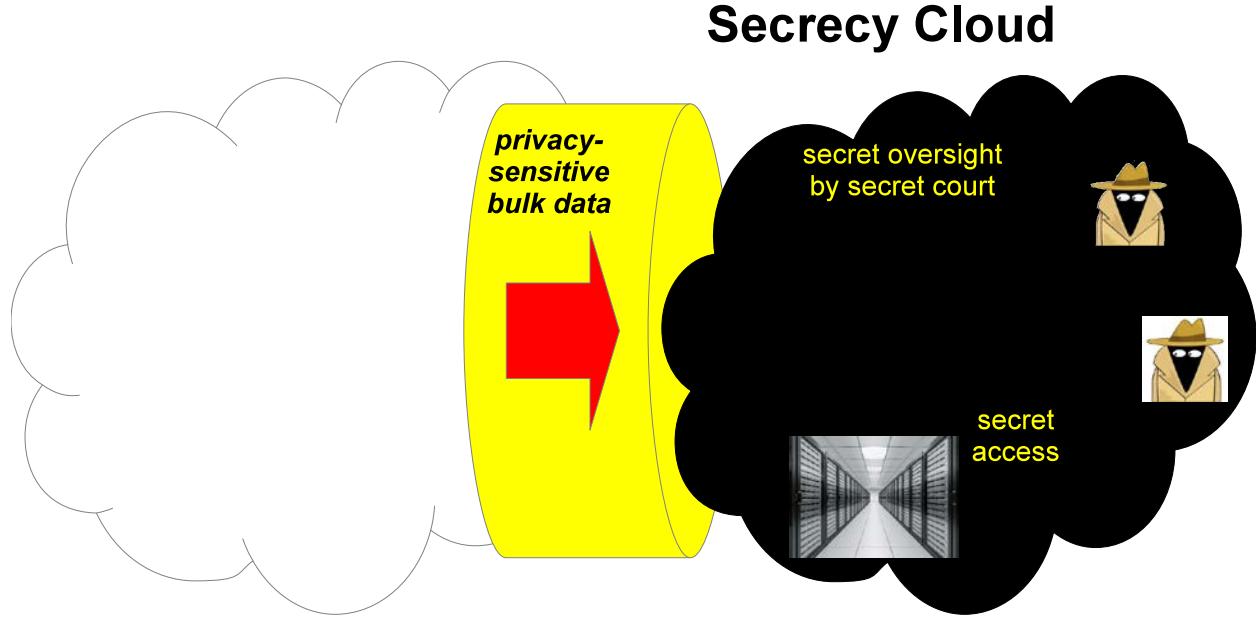


Figure 2: What We Have: A Cloud of Secret Mass Surveillance Processes

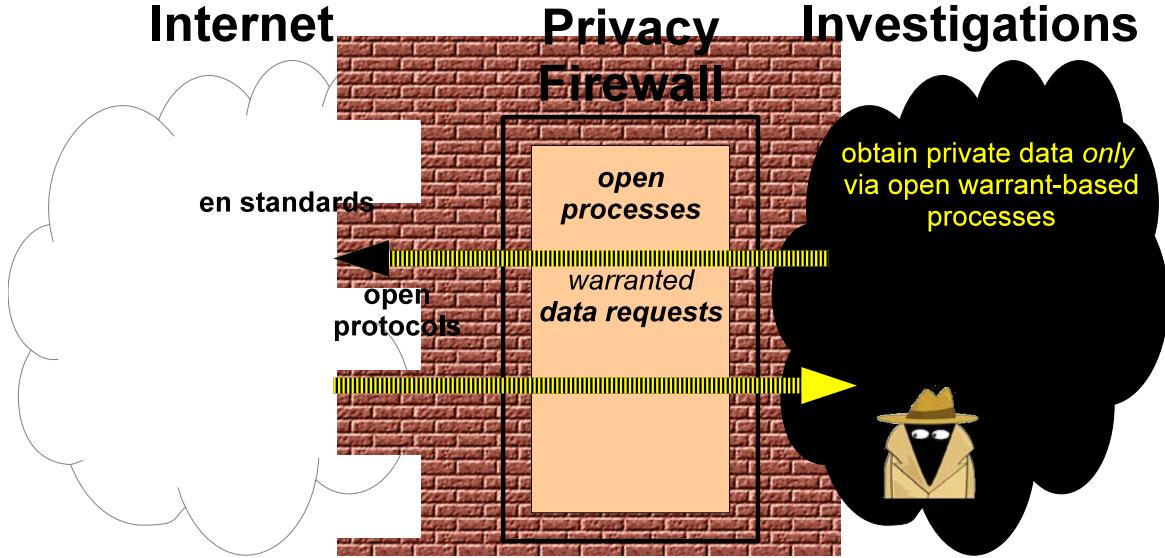


Figure 3: What We Require: Open Warrant-Based Processes for Lawful Electronic Surveillance, Creating a “Privacy Firewall”

Since beginning the PriFi project, we have expanded this work to include an investigation of *contact chaining*, which is also a standard tool of law enforcement and intelligence. The goal is to use the topology of a communication graph (*e.g.*, a phone-call graph, email graph, or social network) to identify associates (or “contacts”) of lawfully targeted users. Agencies can then investigate those associates to determine whether they deserve further attention. It is useful to consider both direct contacts, *i.e.*, users who are neighbors in the communication graph, and extended contacts, *i.e.*, users who are at distance k in the communication graph, for an appropriate constant k . Without mechanisms to limit an investigation’s scope, contact chaining in a mass-communication network can sweep in a huge number of untargeted users.

We present an accountable contact-chaining protocol that bounds the scope of the search, uses encryption to protect untargeted users, and is efficient, with time and communication complexity linear in the size of the output. Experiments show that a three-hop, privacy-preserving graph traversal producing 27,000 ciphertexts can be done in under two minutes.

For a more detailed account, please refer to [10].

5.0 CONCLUSIONS

PriFi is a very promising approach to low-latency, tracking-resistant, local-area networking and could be a critical component of a next-generation communications infrastructure that supports both personal privacy and national security. Preliminary experiments indicate that its performance is good enough to be used in real-life applications, with reasonably high throughput. PriFi’s novel client/relay/server architecture and demanding use of DC-nets are interesting system features.

6.0 REFERENCES

- [1] R. Dingledine, N. Mathewson, and P. Syverson. **Tor: The Second-Generation Onion Router.** 13th USENIX Security Symposium, 2004.
- [2] D. Wolinsky, H. Corrigan-Gibbs, B. Ford, and A. Johnson. **Dissent in Numbers: Making Strong Anonymity Scale.** 10th USENIX Symposium on Operating Systems Design and Implementation, 2012.
- [3] L. Barman, M. Zamani, I. Dacosta, J. Feigenbaum, B. Ford, J.-P. Hubaux, and D. Wolinsky. **PriFi: A Low-Latency and Tracking-Resistant Protocol for Local-Area Anonymous Communication.** 15th ACM Workshop on Privacy in the Electronic Society, 2016.
- [4] J. Weber. **Analysis of Scheduling Algorithms for PriFi.** Thesis: Master in Communication Systems, EPFL, 2017.
- [5] **Analysis of the PriFi Protocol.** Unpublished manuscript, 2017.
URL: <http://www.cs.yale.edu/homes/jf/xxAnalysisxx.pdf>. Accessed June 29, 2017.
- [6] **Shadow: Real Applications, Simulated Networks.** URL: <http://shadow.github.io/>. Accessed June 29, 2017.
- [7] A. Johnson, R. Jansen, A. Jaggard, J. Feigenbaum, and P. Syverson. **Avoiding the Man on the Wire: Improving Tor's Security with Trust-Aware Path Selection.** 24th Symposium on Network and Distributed System Security, 2017.
- [8] E. Syta, P. Jovanovic, E. Kogias, N. Gailly, L. Gasser, I. Khoffi, M. Fischer, and B. Ford. **Scalable Bias-Resistant Distributed Randomness.** 38th IEEE Symposium on Security and Privacy, 2017.
- [9] A. Segal, B. Ford, and J. Feigenbaum. **Catching Bandits and Only Bandits: Privacy-Preserving Intersection Warrants for Lawful Surveillance.** 4th USENIX Workshop on Free and Open Communications on the Internet, 2014.
- [10] A. Segal, J. Feigenbaum, and B. Ford. **Privacy-Preserving Lawful Contact Chaining.** 15th ACM Workshop on Privacy in the Electronic Society, 2016.

7.0 APPENDIX AND BIBLIOGRAPHY

Published Papers, MS Thesis, and Unpublished Manuscript (Each Included Below)

Analysis of the PriFi Protocol. Unpublished manuscript, 2017.

URL: <http://www.cs.yale.edu/homes/jf/xxAnalysisxx.pdf>. Accessed June 29, 2017.

L. Barman, M. Zamani, I. Dacosta, J. Feigenbaum, B. Ford, J.-P. Hubaux, and D. Wolinsky.

PriFi: A Low-Latency and Tracking-Resistant Protocol for Local-Area Anonymous Communication. 15th ACM Workshop on Privacy in the Electronic Society, 2016.

A. Johnson, R. Jansen, A. Jaggard, J. Feigenbaum, and P. Syverson. **Avoiding the Man on the Wire: Improving Tor's Security with Trust-Aware Path Selection.** 24th Symposium on Network and Distributed System Security, 2017.

A. Segal, J. Feigenbaum, and B. Ford. **Privacy-Preserving Lawful Contact Chaining.** 15th ACM Workshop on Privacy in the Electronic Society, 2016.

E. Syta, P. Jovanovic, E. Kogias, N. Gailly, L. Gasser, I. Khoffi, M. Fischer, and B. Ford. **Scalable Bias-Resistant Distributed Randomness.** 38th IEEE Symposium on Security and Privacy, 2017.

J. Weber. **Analysis of Scheduling Algorithms for PriFi.** Thesis: Master in Communication Systems, EPFL, 2017.

LIST OF ACRONYMS

BFT	Byzantine Fault Tolerance
DARPA	Defense Advanced Research Projects Agency
DC-nets	Dining-Cryptographers Networks
EPFL	Ecole Polytechnique Federale de Lausanne
UT	University of Texas
VPN	Virtual Private Network